

In the Claims

1. (Previously presented): A system for storing checkpoint state information, comprising:
 - a network interface to an external network; and
 - a persistent memory unit coupled to the network interface, wherein:
 - the persistent memory unit is configured
 - with byte-level memory access granularity to receive the checkpoint data via a remote direct memory write command from a primary process through the network interface,
 - with meta-data regarding the contents and layout of memory regions within the persistent memory unit, and
 - to provide access to the meta-data and the checkpoint data via a remote direct memory read command from the backup process through the network interface; and
 - the backup process provides recovery capability in the event of a failure of the primary process.
-
2. (Previously presented): The system of Claim 1, further comprising:
 - a persistent memory manager configured to provide address context information to the network interface and to keep the meta-data on the persistent memory unit consistent with the checkpoint data stored on the persistent memory unit.

 3. (Previously presented): The system of Claim 1, wherein the persistent memory unit is configured to provide remote direct memory read access to the checkpoint data to another processor, and the backup process is executed by the other processor.

 4. (Previously presented): The system of Claim 1, wherein the persistent memory unit provides the checkpoint data through remote direct memory reads by the backup process after the primary process fails.

5. (Previously presented): The system of Claim 1, wherein the persistent memory unit is configured to store multiple sets of checkpoint data through remote direct memory writes sent from the processor at successive time intervals.
6. (Previously presented): The system of Claim 5, wherein the persistent memory unit provides the multiple sets of checkpoint data through remote direct memory reads upon request by the backup process at one time.
7. (Previously presented): The system of Claim 1, wherein the primary process remote direct memory writes the checkpoint data to the persistent memory unit independently from the backup process.
8. (Original): The system of Claim 1, wherein the persistent memory unit is configured as part of a remote direct memory access-enabled system area network.
9. (Original): The system of Claim 1, wherein the persistent memory unit is configured with address protection and translation tables to authenticate requests from remote processors, and to provide access information to authenticated remote processors.
10. (Previously presented): A method for recovering the operational state of a primary process, comprising:
- mapping virtual addresses of a persistent memory unit to physical addresses of the persistent memory unit, wherein the persistent memory unit is addressable at byte-level granularity;
- remote direct memory writing checkpoint data regarding the operational state of the primary process to the persistent memory unit;
- remote direct memory reading the checkpoint data from the persistent memory unit via a remote direct memory read command from the backup process;
- storing access information to the physical addresses of the checkpoint data in the persistent memory unit when the primary process opens a memory region for the checkpoint data; and
- providing the access information to subsequent requestors of the checkpoint data.

11. (Original): The method of Claim 10, further comprising:
providing context information regarding the addresses to the primary process and the
backup process.
12. (Previously presented): The method of Claim 10, further comprising:
remote direct memory reading the checkpoint data by the backup process upon failure
of the primary process.
13. (Original): The method of Claim 10, further comprising:
overwriting the checkpoint data with current checkpoint data.
14. (Previously presented): The method of Claim 10, further comprising:
appending updated checkpoint data to at least one previous set of the checkpoint data.
15. (Previously presented): The method of Claim 14, further comprising:
clearing the multiple sets of checkpoint data.
16. (Previously presented): The method of Claim 10, further comprising:
allowing the backup process to remote direct memory read previously unread portions
of the checkpoint data upon failure of the primary process; and
resuming functions performed by the primary process with the backup process.
17. (Canceled)
18. (Currently amended): The method of Claim 10 17, further comprising:
establishing a connection to a process requesting access to the checkpoint data; and
binding the access information to the connection.
19. (Currently amended): The method of Claim 10 17, further comprising:
verifying authentication information from the subsequent requestors.

20. (Original): The method of Claim 10, further comprising:
authenticating a persistent memory manager during initialization of address protection
and translation tables on the persistent memory unit.
21. (Previously presented): A computer product, comprising:
computer executable instructions embodied in a computer readable medium and
operable to:
allow remote direct memory access to a persistent memory unit from a remote
processor via a network, wherein the remote direct memory access
references a persistent memory virtual address;
store checkpoint data from a primary process;
authenticate requests from remote processors, and provide access information
to authenticated remote processors based on address protection and
translation tables in the persistent memory unit;
translate the virtual address to a physical address in the persistent memory unit,
wherein the persistent memory unit is addressable at byte-level
granularity; and
allow access to the checkpoint data for use in a backup process.
22. (Previously presented): The computer product of Claim 21, further comprising:
computer executable instructions embodied in a computer readable medium and
operable to:
allow the processor to access address context information.
23. (Previously presented): The computer product of Claim 21, further comprising:
computer executable instructions embodied in a computer readable medium and
operable to:
store multiple updates to the checkpoint data sent at successive time intervals.

24. (Currently amended) The computer product of Claim 21, further comprising:
computer executable instructions operable to:
allow the backup process to access the multiple sets of the checkpoint data at one time
25. (Original): The computer product of Claim 21, wherein the persistent memory is configured as part of a remote direct memory access-enabled system area network.
26. (Previously presented): An apparatus comprising:
means for communicatively coupling a persistent memory unit to a network that enables remote direct read and write access to a persistent memory unit, wherein the persistent memory unit is addressable at byte-level granularity;
means for receiving access information to physical addresses of checkpoint data in the persistent memory from the persistent memory unit;
means for mapping virtual addresses of the persistent memory unit to physical addresses of the persistent memory unit;
means for receiving the checkpoint data for a primary process in the persistent memory unit via the network; and
means for allowing a backup process to access the checkpoint data via the network.
27. (Previously presented): The apparatus of Claim 26, further comprising:
means for allowing the primary process and the backup process to access context information regarding the addresses.
28. (Previously presented): The apparatus of Claim 26, further comprising:
means for allowing the backup process to access the checkpoint data upon failure of the primary process.

29. (Original): The apparatus of Claim 26, further comprising:
means for creating multiple sets of checkpoint data by appending updated checkpoint
data to at least one previous set of the checkpoint data; and
means for overwriting the checkpoint data with current checkpoint data.
30. (Previously presented) The apparatus of Claim 29, further comprising:
means for periodically accessing at least a portion of the multiple sets of checkpoint
data in the backup process.
31. (Previously presented): The apparatus of Claim 30, further comprising:
means for allowing the backup process to access previously unread portions of the
checkpoint data upon failure of the primary process.
32. (Previously presented): A method for recording the operational state of a primary
process, comprising:
receiving access information to physical addresses of checkpoint data in the persistent
memory from the persistent memory unit; and
accessing checkpoint data regarding the operational state of the primary process in a
persistent memory unit via a remote direct memory access write command,
wherein the persistent memory unit is addressable at byte-level granularity.
33. (Previously presented): The method of Claim 32, further comprising:
overwriting the checkpoint data in the persistent memory unit with current checkpoint
data via a remote direct memory access write command.
34. (Previously presented): The method of Claim 32, further comprising:
appending updated checkpoint data to a previous set of the checkpoint data via a
remote direct memory access write command.

35. (Previously presented): A method for retrieving the operational state of a primary process, comprising:
receiving access information to physical addresses of checkpoint data in the persistent memory from the persistent memory unit; and
transmitting a remote direct memory access read command via network to a remote persistent memory unit from a backup process for the primary process, wherein the persistent memory unit is addressable at byte-level granularity.
36. (Previously presented) The method of Claim 35, further comprising:
periodically transmitting the remote direct memory access read command to retrieve at least a portion of the checkpoint data for the backup process.
37. (Previously presented): The method of Claim 36 35, further comprising:
transmitting the remote direct memory access read command to retrieve previously unread portions of the checkpoint data upon failure of the primary process.